

82771.P277

UNITED STATES PATENT APPLICATION

FOR

METHOD AND APPARATUS FOR VIRTUAL OVERLAY NETWORKS

Inventor(s):

Shantigram Jagannath

Naganand Doraswamy

Andre Fredette

Prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN  
12400 Wilshire Boulevard  
Seventh Floor  
Los Angeles, California 90025

(408) 720-8598

850E90" EHE60T60

BACKGROUND OF THE INVENTION

1. FIELD OF THE INVENTION

The present invention relates to the field of networking and more specifically to virtual overlay networks (VONs) and virtual private networks (VPNs).

2. DESCRIPTION OF THE RELATED ART

Virtual private networks (VPNs) allow users to appear to be on the same private network although there may be many (usually public) networks in between the users. Figure 1A illustrates the logical appearance to users of a virtual private network. Figure 1B illustrates a high level view of the actual network configuration.

Packets destined from one user (say in Chicago in the illustration of Figure 1B) to another user (say in Boston in the illustration of Figure 1B) may be transmitted through an internet service provider (ISP) which supports VPNs. Each site connected to the ISP network advertises to the ISP a set of destinations reachable within the site. The ISP then redistributes this information to all other sites in the set of sites which form the VPN. This process is further described in Heinanen, et al., VPN support with MPLS, Internet Draft, draft-heinanen-mpls-vpn-01.txt, March 1998.

Since the ISP may support multiple VPNs, and since these VPNs may use private address spaces (and, thus the addresses spaces may be non-unique), the routing system within the ISP needs to be able to unambiguously differentiate reachability information (i.e., private address space information)

1 for the various VPNs. Heinanen, et al describes that this may be  
2 accomplished by having the ISP assign each VPN its own VPN identifier  
3 (VPN-ID) and having the routing system use a combination of the VPN-ID  
4 and the reachability information provided by the sites for routing. In such a  
5 system, a single routing system may support multiple VPNs whose address  
6 spaces overlap with each other.

7 Figure 2 illustrates an exemplary prior art routing system using VPN-IDs  
8 and reachability information provided by the sites for routing. As illustrated  
9 by Figure 2, a packet to be routed may include a virtual private network  
10 identifier (VPN-ID) 201, reachability information (e.g., private addressing  
11 information) 202, an internet protocol (IP) header 203 and payload  
12 information 204. A single route table 206 is maintained and is indexed by the  
13 combination of the VPN-ID 201 and the reachability information 202.

14 As is shown in Figure 2, prior art solutions provide a flat address  
15 routing space by simply combining the VPN-ID with the reachability  
16 information provided by the sites.

17 It would be useful to provide more fine-grained control over the routed  
18 topology for individual VPNs.

09109343-063098

SUMMARY OF THE INVENTION

1  
2

3     A method and apparatus for directing messages through a network  
4 wherein multiple tables for directing messages through the network are  
5 maintained and provided. Each table corresponds to a virtual private network  
6 and contains routing information specific to that virtual private network. A  
7 separate routing table is maintained for each virtual private network. In one  
8 embodiment the messages are forwarded using plain IP forwarding, based on  
9 a route table corresponding to the virtual private network. In another  
10 embodiment separate forwarding tables using labels are generated for each  
11 virtual private network. In a third embodiment, a single forwarding table is  
12 utilized where the forwarding table is created based on separate routing  
13 tables for each virtual private network.

860E90" E4E60F60

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1A illustrates a logical representation of a virtual private network.

Figure 1B illustrates a high level conceptual representation of virtual private network.

Figure 2 illustrates a prior art packet / routing table arrangement.

Figure 3A illustrates a first embodiment of a packet / routing table arrangement as may be implemented by the present invention.

Figure 3B illustrates a second embodiment of a packet / routing table arrangement as may be implemented by the present invention.

Figure 3C illustrates a third embodiment of a packet / routing table arrangement as may be implemented by the present invention.

Figure 4 is a high level diagram illustrating a network as may implement the present invention.

For ease of reference, it might be pointed out that reference numerals in all of the accompanying drawings typically are in the form "drawing number" followed by two digits, xx; for example, reference numerals on Figure 1 may be numbered 1xx; on Figure 3, reference numerals may be numbered 3xx. In certain cases, a reference numeral may be introduced on one drawing and the same reference numeral may be utilized on other drawings to refer to the same item.

DETAILED DESCRIPTION OF  
THE EMBODIMENTS THE PRESENT INVENTION

As was discussed above, Virtual private networks (VPNs) allow users to appear to be on the same private network although there may be many (usually public) networks in between the users.

Packets destined from one to another user may be transmitted through an internet service provider (ISP) which supports VPNs. Each site connected to the ISP network advertises to the ISP a set of destinations reachable within the site. The ISP then redistributes this information to all other sites in the set of sites which form the VPN. Since the ISP may support multiple VPNs, and since these VPNs may use private address spaces (and, thus the addresses spaces may be non-unique), the routing system within the ISP needs to be able to unambiguously differentiate reachability information (i.e., private address space information) for the various VPNs.

A similar issue regarding need to unambiguously differentiate reachability information exists with Virtual Overlay Networks (VONs). VONs provide the capability to build logical independent networks over a shared public network infrastructure. VONs are particularly attractive to bandwidth and network infrastructure wholesalers and can also benefit both ISPs and private enterprise networks. VONs allow logical partitioning of networks without building expensive filtering mechanisms. For example, multiple small ISPs could share the same network infrastructure (consisting of, e.g., high bandwidth links and high end router devices) while each ISP could be provisioned to offer specific and tailored services (e.g., real time multicast service) to targeted customers. The concept could also be applied in the

860E90" EHE60T60

1 context of a single ISP when it sells services to different private customers.  
2 Each ISP could have a routed topology that is optimized for its needs – it will  
3 only use those nodes and those links that it requires to provide services. This  
4 logical separation allows a single high bandwidth network infrastructure with  
5 high bandwidth routers to be shared by many small ISPs offering specialized  
6 services. Alternatively it allows a single ISP to partition its network into nodes  
7 and links that are used for specialized services and those that are used to  
8 carry primarily best effort traffic.

9 In the present invention such logically separated routed topologies are  
10 maintained for each VPN. A packet belonging to a VPN is identified by its  
11 VPN-ID. The VPN-ID is placed in the label field as defined by the Multi-  
12 protocol label switching standard, see, Callon et al., A Framework for  
13 Multiprotocol Label Switching, draft-ietf-mpls-framework-02.txt, November,  
14 1997. (Callon et.al). In one embodiment, the VPN-ID is not used for  
15 forwarding, but merely identifies a routing table belonging to a particular VPN.  
16 In this embodiment the packet is forwarded by doing a standard IP  
17 destination address look-up on the table identified by the VPN-ID. In another  
18 embodiment, the VPN-ID identifies an MPLS forwarding table corresponding  
19 to the VPN where the MPLS forwarding table is built based on the routing  
20 table corresponding to the VPN. In a third embodiment, the VPN-ID is a part  
21 of the MPLS forwarding label. A single MPLS forwarding table is built based  
22 on a separate route table for each VPN and the forwarding is done by looking  
23 up the MPLS label (comprising of the VPN-ID part and a forwarding label  
24 part) in the forwarding table.

1 This approach of providing a logically separated routed topology for each  
2 VPN offers significant advantages over prior art approaches. Utilizing this  
3 approach, an ISP may, for example:

1. choose which links and nodes are in a given VPN;
2. assign a given link different administrative weights for different VPNs; and
3. allocate different service levels/guarantees for different VPNs or provisions the service levels and guarantees differently.
4. use different routing protocols for the different VPNs
5. completely isolate the traffic of one VPN from another.

Multi Protocol Label Switching (MPLS) is used on the data plane in certain embodiments of the present invention. MPLS is described in greater detail in Callon et al.,. MPLS is intended to simplify the forwarding function of routing devices by introducing a connection-oriented mechanism inside the otherwise connectionless IP technology. A label switched path (LSP) is set up for each route. Edge routers analyze the traditional IP header (such as IP header 203) to decide which LSP to use and add a corresponding label switched path identifier in the form of a label (such as is show in Figure 3A as VPN-ID 201, in Figure 3B as VPN-ID 201 and forwarding label 302 and in Figure 3C as VPN-ID/forwarding label 311).

21 As will be described, MPLS may be used to facilitate implementation of  
22 logically separated VPNs.

Figure 4 provides a high level overview of a network as may implement the present invention. An edge router (such as router 401, 402 or 403) which resides at the enterprise or alternatively, at the ISP's site. The edge router 401-3 classifies packets onto a given VPN. The packet-to-VPN classification



1 may be based on standard filtering techniques (e.g., input port and IP header  
2 mask). The edge router 401-403 then applies a VPN-specific label to the  
3 packet so that it can be routed by the backbone routers 411-413 in the wide  
4 area network cloud.

5 Three alternative approaches for providing logically separated routed  
6 topologies are described in connection with Figures 3A-C.

7 Turning first to Figure 3A, the label (e.g., VPN-ID 201) is used to  
8 identify a routing table 304 or 305. The packet is then routed based on the  
9 reachability information in the IP header 203. In this embodiment, no label  
10 distribution protocol (e.g. MPLS) is required.

11 Figure 3B illustrates an approach utilizing a label stack comprising the  
12 VPN-ID 201 and a forwarding label 302. In this embodiment, the VPN-ID  
13 201 indicates an MPLS forwarding table 308, 309 which corresponds to the  
14 VPN-ID. The forwarding label 302 provides MPLS label switching  
15 information. By utilizing this embodiment, multiple instances of the standard  
16 MPLS distribution protocol can be utilized.

17 Finally, turning to Figure 3C, each router is allowed to manage its own  
18 MPLS flat label space. Each router is responsible for keeping track of which  
19 VPN each label refers to and for routing them accordingly. The router locally  
20 builds the labels based on its route tables and VPN-ID information and stores  
21 them in a single MPLS forwarding table 312. In this embodiment, the router  
22 still maintains separate route tables for each VPN. In the described  
23 embodiment, as illustrated by Figure 3C, this is accomplished by extending  
24 the label distribution protocol to carry a VPN-ID with the forwarding label as a  
25 combined VPN-ID/forwarding label 311.

26

30

09109343-063098

1

2

3

4

---

Thus, what has been disclosed a method and apparatus for  
maintaining logically separate routing topologies based on virtual private  
networks.

060E90" E4E60T60